WHAT IS CLAIMED IS:

1.   A method of providing highly-available database clusters which replicate at least one data file, the method comprising:

receiving one or more data requests from a client application over a first client connection, the one or more data requests requesting access to data of a first data file;

associating a timestamp with each data request;

forwarding each data request to a first database management system capable of accessing the data of the first data file according to the data request;

forwarding each data request and each associated timestamp to a memory;

forwarding replication data to the memory, wherein the replication data is sufficient to indicate any changes made to the data of the first data file based on the one or more data requests and wherein each replication data includes a timestamp;

comparing the timestamps of one or more of the replication data to the timestamps of one or more of the data requests;

when the one or more of the data requests are determined to be redundant to the one or more of the replication data based on the timestamps, purging one of (a) one or more of the data requests and (b) one or more of the replication data; and

forwarding the other of (a) and (b) to a second database management system capable of accessing data of a second data file.

2.   The method of Claim 1, wherein each timestamp comprises a location ID.

3.   The method of Claim 2, wherein the location ID associated with the one or more data requests corresponds to an indication of a substantially current location in a log file associated with the first database management system.

4.   The method of Claim 2, wherein the location ID associated with the one or more replication data corresponds to the location in a log file associated with the first database management system where the replication data is stored.

5.   The method of Claim 2, wherein the location ID comprises a sequence number and an offset.

6.   The method of Claim 1, wherein at least a portion of the data in the first data file is being replicated to the second data file.

7.   The method of Claim 1, further comprising determining a need to move the first client connection to a second client connection and moving the first client connection to the second client connection.

8.   The method of Claim 7, wherein the move is transparent to the client application such that additional data requests sent by the client are received over the second client connection.

9.   The method of Claim 7, wherein determining the need includes detecting the unavailability of the first database management system.

10.   The method of Claim 9, wherein the unavailability is caused by the failure of the first database management system.

11.   The method of Claim 9, wherein the unavailability is caused by maintenance of one or more portions of the first database management system.

12.   The method of Claim 7, wherein determining the need includes detecting a load on the first database management system.

13.   The method of Claim 7, wherein determining the need includes optimizing a load on the first database management system.

14.   The method of Claim 7, wherein determining the need includes detecting a load on the second database management system.

15.   The method of Claim 7, wherein determining the need includes accessing system configurations of the first database management system.

16.   The method of Claim 15, wherein the system configurations include at least one of hardware and software configurations.

17. A method of performing replication in a database cluster having client connection failover, the method comprising:

monitoring a first system with a second system wherein at least a portion of a first data file of the first system is replicated in a second data file of the second system;

determining a need to move communication over a first connection between a client application and the first data file through the first system to a second connection between the client application and the second data file through the second system;

moving the communication to the second connection in a manner allowing the client application to not fail.

18. The method of Claim 17, wherein the manner includes sending keep-alive messages to the client application from the second system.

19. The method of Claim 17, wherein the determining the need includes detecting the unavailability of the first system.

20. The method of Claim 19, wherein the unavailability is caused by the failure of the first database management system.

21. The method of Claim 19, wherein the unavailability is caused by maintenance of one or more portions of the first system.

22. The method of Claim 17, wherein determining the need includes detecting a load on the first system.

23. The method of Claim 17, wherein determining the need includes optimizing loads between the first and second systems.

24. The method of Claim 17, wherein determining the need includes detecting a load on the second system.

25. The method of Claim 17, wherein determining the need includes determining system configurations of the first system.

26. The method of Claim 25, wherein the system configurations include at least one of hardware and software configurations.